

The Maintenance Constraint: How Resource Boundaries Shape Cognitive Availability

Lluis Eriksson
Independent Researcher
lluiseriksson@gmail.com

December 2025

Abstract

Cognitive systems are resource-limited, but “resource limitation” is often invoked without distinguishing one-shot costs (forming a representation) from sustained costs (keeping it usable under noise). We argue that *availability* of internal state features for control, integration, and report is constrained by their *maintainability* under finite budgets.

As a technical anchor, we cite a companion technical preprint deriving an operational inequality in explicit thermodynamic control models, showing that incremental maintenance power can have a non-arbitrary lower bound tied to dynamical fragility. This motivates an *operational cut*: a feasibility boundary separating maintainable from unmaintainable state features.

We develop an auditable bridge argument (maintainability \rightarrow stability \rightarrow availability), propose a neutrality-friendly principle of *maintenance-feasibility bias*, and show how it can be incorporated into active inference (Free-Energy Principle) as a maintenance penalty or constraint, while aligning secondarily with Global Workspace accounts of access stability. We address common objections and offer falsifiable predictions for synthetic agents and neuromorphic systems, plus an explicitly exploratory psychophysics subsection framed in terms of reportability and stability rather than phenomenology. We do not propose collapse mechanisms, do not derive the Born rule, and make no claims about phenomenological consciousness.

1 Introduction

Many debates in philosophy of mind and cognitive neuroscience presuppose that internal variables can be treated as stable quantities available for downstream computation, control, and report. Yet biological and artificial agents operate under noise, drift, coupling, and finite budgets (metabolic power, bandwidth, control precision, compute). This raises a neglected question: *which internal state features are feasible to keep available over task-relevant timescales?*

A common move is to invoke “resource limits” in broad terms. While true, this can be too coarse to constrain architectures. The aim of this paper is to sharpen the idea into an operational constraint: availability is limited by maintainability. Some state features may be representable in principle but too fragile or too costly to keep stably available under realistic budgets. This shapes which representations dominate behavior and report.

States vs. state features. We often speak of *state features* rather than exact microstates. In realistic control and cognitive settings, what must be stabilized is typically a task-relevant feature (a coarse-graining, a decoded variable, a functional manifold, a correlation pattern), not an entire microscopic configuration.

Two methodological commitments guide the paper:

1. We treat “availability” as a functional notion (access for downstream use), not as a claim about phenomenology. This targets access-level cognition and report, not the hard problem [2, 3, 4].
2. We keep the physics anchor separate from philosophy-of-mind claims. The anchor supplies an example of a nontrivial lower bound on maintenance cost in an explicit model; it does not imply that brains implement quantum coherence, nor that consciousness depends on quantum effects.

The paper is structured as follows. Section 2 surveys existing approaches to physical constraints on cognition. Section 3 summarizes the technical anchor and defines the operational cut. Section 4 develops the bridge argument (maintainability \rightarrow stability \rightarrow availability) and introduces maintenance-feasibility bias, including operational proxies. Section 5 integrates the constraint into active inference (FEP) and aligns it with Global Workspace accounts. Section 6 addresses objections. Section 7 proposes falsifiable implications for synthetic and neuromorphic systems, plus an explicitly exploratory psychophysics subsection framed in terms of reportability and stability. Section 8 concludes.

Non-claims (explicit). We do not propose collapse mechanisms, do not derive the Born rule, and make no claims about phenomenological consciousness. The term “cut” is used operationally: a feasibility boundary under resource constraints, not an ontological boundary in nature.

2 Existing approaches to physical constraints on cognition

The idea that cognition is physically constrained is not new. What is often missing is a principled separation between constraints on *computation/representation* and constraints on *maintenance of state features over time under noise and coupling*. We briefly situate the present proposal.

2.1 Computational and informational constraints

Several traditions emphasize computational limits: bounded rationality, limited memory, limited attention, and limited communication [11]. In cognitive science, Marr’s levels distinguish computational, algorithmic, and implementational analyses [10]. The implementational level motivates energetic and physical constraints, but implementational constraints are often treated as background rather than as central architecture-shaping principles.

Related work in information theory and neuroscience analyzes channel capacity, coding efficiency, and limited bandwidth. While powerful, such approaches often constrain what can be encoded or transmitted, not necessarily what can be kept stably available under ongoing perturbations.

2.2 Energetic constraints in neural systems

Neuroenergetics provides direct evidence that maintenance is costly: sustaining membrane potentials against leakage and supporting spiking and synaptic transmission consume substantial metabolic power [13, 14]. This motivates the general expectation that stable internal dynamics are not free.

However, energetic discussions in neuroscience often remain at the level of total budget or average cost per spike. What is less developed is a bridge from energetic limits to constraints on which internal variables can remain stable and available as cognitive representations under specific noise profiles and couplings.

2.3 Resource rationality and the missing dynamical dimension

Recent work on resource rationality emphasizes that agents optimize objectives subject to resource costs [12]. This framework shares our emphasis on constraints but differs in focus:

- **Resource rationality** typically concerns computational costs: time, memory, samples, or algorithmic complexity required to reach a decision.
- **Maintenance feasibility** concerns dynamical costs: the ongoing expenditure required to keep internal state features within functional bounds under noise and coupling.

These are complementary rather than competing. A representation may be computationally cheap to derive but dynamically expensive to hold (or vice versa). A related distinction exists in control theory between reachability (can the system reach a target?) and stabilizability (can it stay there under perturbations?). Our contribution is to import this distinction into cognitive architecture via the maintainability–availability bridge.

2.4 What the operational cut adds

Existing approaches treat resource limits as background constraints on computation or total energy. What is missing is a principle that specifically targets which internal state features can remain stably available under noise and coupling. The operational cut provides this by linking maintainability to stability and then to availability. Unlike generic bounded rationality, it yields testable predictions about how resource scarcity shapes representational forms (compression, buffering, modularization) rather than merely bounding capacity.

3 Technical anchor: a maintenance inequality and an operational cut

This section summarizes the technical anchor at a level sufficient for the philosophical argument, without requiring the reader to follow the full physics proof.

3.1 The anchor inequality (summary)

In an explicit thermodynamic control model (finite-dimensional system, Markovian uncontrolled dynamics, and a work reservoir/battery interacting via energy-conserving global unitaries with a thermal environment at temperature T), one can define an incremental maintenance power $P_{\text{extra}}(\rho)$ that isolates the steady cost of stabilizing certain state features at fixed populations. A central result is a lower bound of the form

$$P_{\text{extra}}(\rho) \geq k_{\text{B}}T \dot{C}_{\text{loss}}(\rho),$$

where $C(\rho)$ is a relative-entropy-based coherence functional and $\dot{C}_{\text{loss}}(\rho)$ is the instantaneous loss rate of that coherence under the uncontrolled dynamics. The important feature for our purposes is structural: maintenance has a non-arbitrary lower bound tied to dynamical fragility. This can be viewed as a paradigmatic instance of a maintenance constraint.

We emphasize: this anchor is used as a benchmark case study of maintainability costs, not as a claim about neural implementation. The philosophical argument does not require that the brain maintains quantum coherence. The relevant import is the existence of principled maintenance lower bounds in explicit physical models.

3.2 The operational cut

Given a finite available budget P_{avail} , we can define a feasibility boundary:

Operational cut: the boundary between state features that can be maintained with $P_{\text{extra}}(\rho) \leq P_{\text{avail}}$ and those that cannot.

This notion of cut is operational and agent-relative: it depends on available resources and the uncontrolled dynamics. It is not an ontological division in nature. Its cognitive role emerges when we connect maintainability to stability and availability.

3.3 Why this matters philosophically

The anchor shows how a maintainability question becomes a budget-feasibility question: some state features are too expensive to keep from degrading. This supports a general methodological point: if availability requires stability, and stability requires maintainability, then resource boundaries constrain which internal state features can function as stable cognitive representations.

3.4 Methodological status of the technical anchor

Our use of the anchor result is methodological rather than neurobiological. The anchor provides an existence proof that, in explicit physical control models, holding certain state features against uncontrolled dynamics entails non-arbitrary lower bounds tied to dynamical fragility. We do not infer that brains instantiate the anchor’s microphysics. Rather, we use it to motivate a general research stance: in any noisy physical substrate, availability of fine-grained internal variables over task-relevant horizons may require ongoing stabilization whose cost is constrained by the substrate’s loss rates and coupling structure. The philosophical claim defended here concerns the architecture-shaping role of such feasibility boundaries, not the universality of any specific inequality.

4 From maintainability to cognitive availability

This section states the central bridge claim as an explicit, auditable argument. The claim is not that maintainability is consciousness. Rather, maintainability constrains which internal state features can stably participate in cognition in the sense of being available for control, integration, and report.

4.1 Three notions: maintainability, stability, availability

We distinguish three notions that are frequently conflated in discussions of physical constraints on cognition.

Maintainability (control-feasibility). A target internal state feature is maintainable relative to a system’s uncontrolled dynamics and control primitives if there exists an admissible strategy that keeps the system in (or near) a functionally acceptable set for the requisite duration, subject to the agent’s resource budget. In the anchor model, maintainability is constrained by a lower bound on incremental maintenance power.

Stability (dynamical robustness). A state feature is stable over a time horizon if, under actual uncontrolled perturbations and couplings, it does not drift rapidly in ways that destroy its functional role. Stability concerns trajectory timescales, not merely instantaneous existence of an encoding.

Availability (cognitive-functional). A state feature is available if it can be reliably used by the system for downstream tasks: guiding action, updating beliefs, broadcasting to other subsystems, or supporting behavioral report. Following the access/phenomenal distinction [2, 5], we treat availability as separable from phenomenology. The present paper concerns availability and control, not qualia.

Maintenance resources (non-quantum reading). By maintenance resources we mean measurable ongoing costs such as metabolic power, control effort, communication bandwidth, precision/feedback gain, or other continuously expended quantities required to keep relevant internal variables within functional bounds.

4.2 Operational proxies for stability and availability

To keep the proposal empirically anchored, we note two families of operational proxies.

Stability proxies. Over a horizon T , stability can be operationalized as (i) a mean deviation from a functional set \mathcal{K} , (ii) a dwell-time in \mathcal{K} before exit, or (iii) a decay time of an autocorrelation of a task-relevant internal variable. These proxies quantify whether a putative representation persists under the relevant uncontrolled perturbations.

Availability proxies. Availability can be operationalized by downstream readout performance at latency τ : how reliably another subsystem can use the variable for control or report. Concretely, one may measure (i) decoding accuracy under noise, (ii) control performance degradation when the variable is perturbed, or (iii) mutual information between the internal variable and downstream actions within the relevant time window. The maintenance constraint predicts systematic trade-offs between these proxies and measured maintenance expenditure under budget manipulations.

4.3 The dependency chain: explicit premises and conclusion

We present the bridge as premises and a conclusion.

P1 (Availability requires task-relevant persistence). For an internal state feature to be available for control, integration, or report, it must persist on timescales comparable to downstream latencies (integration windows, action selection, working-memory update). A microstate overwritten by uncontrolled dynamics before it can influence downstream processes is not available in the relevant sense.

P2 (Persistence under noise often requires maintenance). In realistic agents, uncontrolled dynamics (noise, coupling, drift) tend to degrade fine-grained internal state features toward default configurations. Unless the target lies within a naturally stable manifold (e.g., a deep attractor basin) or is protected by architecture (buffering, modularization), persistence typically requires maintenance (feedback stabilization, metabolic expenditure, error correction, or active control).

P3 (Maintenance can have unavoidable lower bounds in explicit physical models). Under explicit thermodynamic bookkeeping, sustaining deviations from default relaxation can impose non-zero lower bounds on ongoing resource use. The technical anchor provides a paradigmatic case (quantum coherence) used here as an operational benchmark, not as a claim about neural implementation. Specifically, in the anchor model one has $P_{\text{extra}}(\rho) \geq k_{\text{B}}T \dot{C}_{\text{loss}}(\rho)$, linking incremental maintenance power to intrinsic dynamical fragility.

C (Availability is resource-constrained). From P1–P3: the set of state features available to an agent is constrained by the agent’s resource budget. Some representational candidates may exist as encodings in principle but be unavailable in practice because they are too costly to maintain under the actual noise and coupling conditions.

Maintenance-feasibility bias: when multiple internal encodings can serve comparable functional roles, finite-budget agents tend to adopt encodings whose stability is achievable within available maintenance resources.

[What the conclusion does not say] The conclusion is not that cognition is nothing but maintenance cost, nor that phenomenology is determined by the bound. The conclusion is a constraint on availability and stability of internal variables under finite budgets.

4.4 Loading vs. holding: work vs. power

A central operational distinction is between loading a representation (one-shot shaping cost to enter a target state) and holding a representation (sustained maintenance cost to keep it from drifting under perturbations). This disentangles:

1. The system can encode X (existence of an encoding).
2. The system can stably use X (availability under time and noise).

An encoding may be cheap to load but expensive to hold, or vice versa. This predicts dissociations: transient high-fidelity microstates that are not available for report, versus lower-fidelity states that dominate behavior because they persist.

4.5 Passive stability as a special case

A natural objection is that not all stable representations require active maintenance: strong attractors, passive memory traces, and structural constraints can yield persistence without ongoing work.

Reply: passive stability is a special case where the uncontrolled dynamics already keep the state near target, so incremental maintenance cost can be negligible. In the anchor model, this corresponds to vanishing loss rate for the relevant state feature, making the bound trivial. The maintenance-feasibility bias reduces to the correct special case: availability is unconstrained by maintenance when dynamics are already aligned.

4.6 Preview: maintainability as a regularizer in active inference

Under active inference, we can treat maintenance as a constraint or penalty on internal dynamics, yielding a multi-objective trade-off between variational objectives and maintainability costs. Section 5 develops this connection.

5 Connections to cognitive frameworks

We connect the maintenance constraint to two influential frameworks: active inference (FEP) as the primary backbone, and global workspace accounts as a secondary alignment for access and reportability.

5.1 Active inference / Free-Energy Principle with a maintenance penalty

Active inference proposes that agents minimize a variational free-energy functional as a proxy for surprise, under a generative model linking hidden states and observations [8, 9]. As a general strategy, this is compatible with the idea that agents choose internal states and policies to reduce expected surprise.

The maintenance constraint adds an implementational dimension: even if a representational state reduces variational objectives, it may be too fragile to hold under realistic perturbations. This motivates a multi-objective formulation

$$\min_{q, \pi} \left\{ \mathbb{E}_{\pi} [\mathcal{F}_{\text{var}}(q)] + \lambda \mathcal{M}(q, \pi) \right\},$$

where q denotes an internal encoding (e.g., a variational density), π is a policy/control variable, \mathcal{F}_{var} is the variational objective, and \mathcal{M} is a maintenance term representing ongoing implementational costs required to keep q within functional bounds (e.g., precision/feedback gain, bandwidth use, metabolic proxies, or control effort).

Maintenance as an implementational regularizer (link to complexity). In many presentations of variational free energy, a central decomposition distinguishes an accuracy-like term from a complexity-like term (divergence from priors). Independently of the exact decomposition adopted, active inference already treats some representational forms as more costly than others in an informational sense.

The maintenance constraint adds a dynamical dimension: some encodings are not only informationally complex but also dynamically fragile under the agent’s noise profile, thereby requiring sustained stabilization to remain available. This turns stability into an explicit regularizer: fragile encodings are penalized not merely because they diverge from priors but because they are costly to hold against drift.

[Constrained vs. penalized form] Equivalently, one may minimize \mathcal{F}_{var} subject to a hard maintenance budget $\mathcal{M}(q, \pi) \leq M_{\text{max}}$. In that constrained view, M_{max} defines an operational cut: encodings with $\mathcal{M} > M_{\text{max}}$ are excluded from the feasible set, regardless of their variational attractiveness.

[On the status of the maintenance term] In cognitive models, $\mathcal{M}(q, \pi)$ is an implementational cost term (e.g., control effort, precision gain, bandwidth use, metabolically relevant proxies). While such proxies may correlate with physical power in specific implementations, the present argument does not assume a one-to-one identification. The substantive claim is that making \mathcal{M} explicit allows stability/availability constraints to enter normative objectives in a testable way.

5.2 Global workspace and access stability (secondary alignment)

Global Workspace and related models emphasize that conscious access (in the reportability sense) involves global broadcasting and integration across subsystems [7, 6]. Such broadcasting plausibly requires stability over integration windows and communication latencies. The maintenance constraint provides a physically grounded reason why some encodings fail to become globally available: even if representable, they may not be maintainable long enough under budgets to support stable broadcasting.

This is not a reduction of workspace to energy expenditure. It is a constraint: architectures requiring long-range, high-precision coordination may incur maintenance costs that exceed available budgets, thereby limiting stable access. This predicts that architectures (biological or synthetic) will favor buffered, modular, low-interference representations for stable broadcast.

5.3 Attention as resource allocation

Attention is often described as limited resource allocation (precision weighting, gain modulation, selection). The maintenance-feasibility bias adds a physical interpretation: attention can be seen as allocating maintenance resources to keep some variables within functional bounds while allowing others to decay. This is compatible with FEP-style precision control and workspace-style selection mechanisms.

5.4 Relation to IIT (brief, non-committal)

Integrated Information Theory (IIT) proposes that consciousness corresponds to integrated information in a system’s causal structure [15]. While IIT focuses on what a system represents (information structure), our framework concerns whether representational candidates can be stably maintained under finite budgets. These are orthogonal: a high-integration configuration might be unmaintainable under strong noise or limited budgets, while a low-integration configuration might be highly stable. A possible research direction is to study whether maintenance costs constrain which information-theoretic structures are realizable in practice.

6 Objections and replies

We address predictable objections.

6.1 Objection 1: “This merely rephrases computational limits”

Reply: computational limits constrain what can be computed or represented. The maintenance constraint targets a different dimension: which internal variables can be kept stably available over time under noise and coupling. It yields operational predictions about stability under budget manipulations and distinguishes loading from holding (work vs. power). This is more specific than generic bounded rationality.

6.2 Objection 2: “The constraint is trivial in classical systems”

Objection: If the brain is classical, the quantum anchor is irrelevant, and the constraint reduces to the truism that systems have limited energy.

Reply: The objection conflates two claims: (a) maintenance has costs (trivial), and (b) maintenance costs can have principled lower bounds tied to dynamical fragility (non-trivial). The anchor demonstrates (b) in an explicit model: the bound is not an engineering inefficiency but a constraint tied to loss rates under uncontrolled dynamics. The philosophical import is structural: if we want to know which representations are maintainable, we must characterize the relationship between fragility and the ongoing cost of stabilization.

Moreover, the maintenance-feasibility bias yields falsifiable predictions in purely classical systems (Predictions 1–3): budget manipulations and architectural collars should change stability and distal influence in measurable ways. The proposal is therefore not vacuous in the classical limit.

6.3 Objection 3: “This does not address phenomenology”

Reply: correct. The paper concerns availability, stability, control, and reportability. We explicitly avoid claims about qualia and do not attempt to solve the hard problem. Nevertheless, many empirical measures in consciousness science track access/reportability rather than phenomenology directly. A constraint on stability and availability is therefore relevant to a significant portion of the empirical and theoretical landscape.

6.4 Objection 4: “Maintainability is not consciousness (thermostat objection)”

Reply: agreed. A thermostat maintains a low-dimensional variable within a narrow range. Our claim is not that maintenance is sufficient for consciousness. The claim is that access-level cognition (integration, flexible control, stable report) plausibly requires maintaining higher-dimensional and longer-range patterns over task-relevant horizons. Such patterns are typically more fragile under noise and coupling and therefore more likely to incur nontrivial maintenance burdens. This is a constraint on availability, not a criterion for phenomenology.

6.5 Objection 5: “The operational cut is model-dependent and therefore not fundamental”

Reply: yes: the cut depends on P_{avail} , the uncontrolled dynamics, and the control primitives. This is a feature rather than a bug. Different systems, tasks, and environments produce different feasibility boundaries. This yields empirical content: we can move the cut by changing budgets, noise, or coupling structure, and test the predicted changes in representational stability and availability.

6.6 Objection 6: “Availability entails stability by definition”

Reply: even if one grants that availability is conceptually tied to task-relevant persistence, the paper’s substantive claims are not merely definitional. The nontrivial steps are that (i) persistence under realistic noise often requires maintenance, (ii) maintenance draws on finite budgets, and (iii) these feasibility constraints predict systematic, testable biases in representational form (compression, buffering, modularization) and in architecture-level sensitivity to distal perturbations.

6.7 Objection 7: “This is a slippery slope to reductionism”

Reply: the constraint is a regularizer, not a reductive base. It does not claim that consciousness reduces to power expenditure. It claims that resource feasibility shapes the space of available representations. This is consistent with non-reductive physicalism, emergentism, or functionalism. The constraint is at the level of architecture and availability, not phenomenal content.

7 Testable implications

We propose falsifiable implications with an emphasis on synthetic agents and neuromorphic systems (cleanest control), plus an explicitly exploratory psychophysics subsection framed in terms of reportability and stability.

7.1 Predictions for synthetic agents and neuromorphic systems (primary)

Prediction 1 (budget manipulation \rightarrow representational compression). **System:** an active-inference-inspired agent implemented in silico with an explicit energy/compute budget.

Manipulation: reduce P_{avail} (or an explicit maintenance-cost parameter λ) while holding task demands fixed.

Prediction: the agent shifts toward encodings with lower fragility, expressed as increased compression/sparsity or reduced effective dimensionality, even if this slightly worsens variational fit.

Metric: effective dimension of internal states, stability of belief trajectories, or variance of posterior dynamics under matched noise.

Success criterion: stability proxies improve or remain stable while dimensionality/compression

increases as P_{avail} decreases.

Failure mode: no systematic change in stability/encoding form despite large budget shifts, suggesting other constraints dominate.

Prediction 2 (architectural collars → reduced distal interference). **System:** neuro-morphic network or recurrent architecture with tunable coupling between a local module and a distal perturbation source.

Manipulation: introduce a “collar” by attenuating couplings along the path (reduced gain, reduced bandwidth, modular buffering).

Prediction: local stability measures become less sensitive to distal perturbations, reducing a dynamical influence proxy at fixed horizon.

Metric: mutual information or coherence proxy between distant modules; change in local state trajectories under matched perturbations.

Success criterion: distal perturbation sensitivity decreases with collar strength at matched task performance.

Prediction 3 (loading vs. holding dissociation). **System:** controlled synthetic agent with explicit control cost.

Manipulation: compare two strategies: one optimized for low one-shot loading cost, one optimized for low sustained holding cost.

Prediction: strategies dissociate: low loading cost need not imply low holding cost, and vice versa.

Metric: one-shot energy/compute spent to reach target vs. average maintenance expenditure over a horizon, at matched performance.

7.2 Exploratory psychophysics implications (secondary, higher risk)

We include one cautious, exploratory line relevant to consciousness science without making phenomenological claims.

Prediction 4 (stability-reportability coupling under resource stress). **System:** human perception in tasks with known bistability or integration windows (e.g., binocular rivalry, continuous flash suppression).

Manipulation: conditions plausibly reducing effective maintenance resources or increasing noise (fatigue, attentional load, sensory noise; ethically appropriate manipulations only).

Prediction: decreased stability of reportable content (shorter dominance durations, increased switching, reduced confidence/consistency), consistent with a shifted operational cut for stable availability.

Metric: dominance duration distributions, report consistency, reaction times, confidence ratings.

Caveat: confounds are substantial; this is a directional hypothesis for targeted study.

What would count as falsification? This exploratory prediction would be weakened if manipulations plausibly reducing maintenance resources reliably increase report stability, or if careful modeling shows dominance dynamics are fully explained by feedforward factors with no sensitivity to resource manipulations at fixed stimulus conditions.

7.3 What would count as falsification (overall)?

For maintenance-feasibility bias, falsification would consist of robust counterexamples where:

- reducing explicit maintenance budgets does not reduce stability/availability proxies in controlled synthetic systems, despite increased noise/coupling, and

- architectural collars do not reduce distal influence in settings where they demonstrably reduce coupling strength, and
- no measurable trade-off exists between stability and maintenance expenditure across policies in budgeted agents.

Because the proposal is operational, failures should be interpreted at the level of assumptions (what counts as maintenance cost, which proxy is appropriate, whether the manipulation truly changes maintainability).

8 Conclusion

We argued that maintainability constraints impose a principled boundary on cognitive availability: internal state features must persist over task-relevant timescales to be available for control, integration, and report, and persistence under noise often requires maintenance under finite budgets. A companion technical preprint provides a benchmark instance in which maintenance has non-arbitrary lower bounds tied to dynamical fragility. This motivates an operational cut: a feasibility boundary between maintainable and unmaintainable state features.

We integrated this constraint into active inference as a maintenance penalty, aligned it with global workspace notions of access stability, addressed common objections, and proposed falsifiable implications emphasizing synthetic and neuromorphic systems. The framework is intentionally non-ontological and avoids claims about collapse, the Born rule, or phenomenological consciousness. Its value lies in offering a clear, testable way to connect resource feasibility to representational stability and availability.

References

- [1] L. Eriksson, *Operational Coherence Maintenance and the Quantum–Classical Boundary: Formal Definitions, Falsifiable Protocols, and an Outlook for Cognitive Systems*, preprint (2025).
- [2] N. Block, On a confusion about a function of consciousness, *Behav. Brain Sci.* 18, 227–287 (1995).
- [3] D. J. Chalmers, Facing up to the problem of consciousness, *J. Conscious. Stud.* 2, 200–219 (1995).
- [4] D. J. Chalmers, *The Conscious Mind* (Oxford University Press, 1996).
- [5] S. Dehaene and L. Naccache, Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework, *Cognition* 79, 1–37 (2001).
- [6] S. Dehaene, *Consciousness and the Brain* (Viking, 2014).
- [7] B. J. Baars, *A Cognitive Theory of Consciousness* (Cambridge University Press, 1988).
- [8] K. Friston, The free-energy principle: a unified brain theory?, *Nat. Rev. Neurosci.* 11, 127–138 (2010).
- [9] T. Parr and K. Friston, Generalised free energy and active inference, *Biol. Cybern.* 113, 495–513 (2019).
- [10] D. Marr, *Vision* (W. H. Freeman, 1982).
- [11] H. A. Simon, A behavioral model of rational choice, *Q. J. Econ.* 69, 99–118 (1955).

- [12] F. Lieder and T. L. Griffiths, Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources, *Behav. Brain Sci.* 43, e1 (2020).
- [13] S. B. Laughlin, J. M. de Ruyter van Steveninck, and J. C. Anderson, The metabolic cost of neural information, *Nat. Neurosci.* 1, 36–41 (1998).
- [14] D. Attwell and S. B. Laughlin, An energy budget for signaling in the grey matter of the brain, *J. Cereb. Blood Flow Metab.* 21, 1133–1145 (2001).
- [15] G. Tononi, M. Boly, M. Massimini, and C. Koch, Integrated information theory: from consciousness to its physical substrate, *Nat. Rev. Neurosci.* 17, 450–461 (2016).